CLAIMS

1. A memory hub, comprising:

a local queue adapted to receive local memory responses, and operable to store the local memory responses;

a bypass path adapted to receive downstream memory responses, and operable to pass the downstream memory responses;

a buffered queue coupled to the bypass path and operable to store downstream memory responses;

a multiplexer coupled to the local queue , buffered queue and bypass path, the multiplexer being operable to output responses inform a selected one of the queues or the bypass path responsive to a control signal; and

arbitration control logic coupled to the multiplexer, the arbitration logic operable to develop the control signal to control the selection of responses output by the multiplexer.

2. The memory hub of claim 1 wherein the arbitration control logic develops the control signal to output memory responses stored in the local queue prior to memory responses stored in the buffered queue.

3. The memory hub of claim 1 wherein the arbitration control logic develops the control signal to output memory responses stored in the buffered queue prior to memory responses stored in the local queue.

4. The memory hub of claim 1 wherein the arbitration control logic develops the control signal to provide the memory responses from the local and buffered queues as a function of the age of a memory request associated with each memory response.

5. The memory hub of claim 4 wherein the arbitration control logic assigns a time stamp to each memory request when the request is received by the hub, and wherein the age of each request corresponds to the assigned time stamp.

6. The memory hub of claim 5 wherein the arbitration control logic develops the control signal to alternately output a predetermined number of memory responses stored in the buffered queue and a predetermined number of memory responses stored in the local queue.

7. The memory hub of claim 6 wherein the same predetermined number of responses are alternately output from each queue if such predetermined number or a greater number of responses are stored in each queue.

8. The memory hub of claim 1 wherein each of the local and downstream memory responses comprise data and a header identifying a memory request corresponding to the memory response.

9. A memory hub adapted to receive local memory responses and downstream memory responses, the memory hub operable to store the received memory responses and operable to apply an arbitration algorithm to select the order in which the stored local and downstream memory responses are provided on an uplink output.

10. The memory hub of claim 9 wherein the memory hub further comprises a local queue that stores the local memory responses and a buffered queue that stores the downstream memory responses.

11. The memory hub of claim 10 wherein the memory hub further comprises a multiplexer coupled to the local queue, buffered queue, and bypass path, the multiplexer providing responses from one of the queues or bypass path on an output responsive to a control signal.

12. The memory hub of claim 11 wherein the memory hub further comprises arbitration logic coupled to the queues and the multiplexer, and wherein the arbitration logic applies the control signal to the multiplexer to control which memory responses are provided on the output.

13. The memory hub of claim 12 further including a bypass path coupled to the buffered queue and coupled to the multiplexer, the bypass path adapted to receive the downstream memory responses and operable to provide the responses to the multiplexer and the buffered queue.

14. The memory hub of claim 9 wherein each of the local and downstream memory responses comprise data and a header identifying a memory request corresponding to the memory response.

15. A memory module, comprising:

a plurality of memory devices; and

a memory hub coupled to the memory devices, the memory hub including,

a local queue adapted to receive local memory responses, and operable to store the local memory responses;

a bypass path adapted to receive downstream memory responses, and operable to pass the downstream memory responses;

a buffered queue coupled to the bypass path and operable to store downstream memory responses;

a multiplexer coupled to the local queue, buffered queue and bypass path, and operable to output responses from one of the queues or the bypass path responsive to a control signal; and

arbitration control logic coupled to the multiplexer, the arbitration logic operable to develop the control signal to control the selection of responses output by the multiplexer.

16. The memory module of claim 15 wherein each of the memory devices comprises an SDRAM.

17. The memory module of claim 16 wherein the arbitration control logic develops the control signal to output memory responses stored in the local queue prior to memory responses stored in the buffered queue.

18. The memory module of claim 17 wherein the arbitration control logic develops the control signal to output memory responses stored in the buffered queue prior to memory responses stored in the local queue.

19. The memory module of claim 16 wherein the arbitration control logic develops the control signal to provide the memory responses from the local queue, buffered queue and the bypass path as a function of the age of a memory request associated with each memory response.

20. The memory module of claim 19 wherein the arbitration control logic assigns a time stamp to each memory request when the request is received by the hub, and wherein the age of each request corresponds to the assigned time stamp.

21. The memory module of claim 16 wherein the arbitration control logic develops the control signal to alternately output a predetermined number of memory responses stored in the buffered queue and a predetermined number of memory responses stored in the local queue.

22. The memory module of claim 21 wherein the same predetermined number of responses are alternately output from each queue if such predetermined number or a greater number of responses are stored in each queue.

23. The memory module of claim 15 wherein each of the local and downstream memory responses comprise data and a header identifying a memory request corresponding to the memory response.

24. A memory system, comprising:

a memory hub controller;

a plurality of memory modules, each memory module being coupled to adjacent memory modules through respective high-speed links, at least one of the memory modules being coupled to the memory hub controller through a respective high-speed link, and each memory module comprising:

a plurality of memory devices; and

a memory hub coupled to the memory devices, the memory hub comprising,

a local queue adapted to receive local memory responses, and operable to store the local memory responses;

a bypass path adapted to receive downstream memory responses, and operable to pass the downstream memory responses;

a buffered queue coupled to the bypass path and operable to store downstream memory responses;

a multiplexer coupled to the local queue, the buffered queue and the bypass path, and operable to output responses from one of the queues or the bypass path responsive to a control signal; and

arbitration control logic coupled to the multiplexer, the arbitration logic operable to develop the control signal to control the selection of responses output by the multiplexer.

25. The memory system of claim 24 wherein each of the high-speed links comprises an optical communications link.

26. The memory system of claim 24 wherein at least some of the memory devices comprise SDRAMs.

27. The memory system of claim 24 wherein the arbitration control logic develops the control signal to output memory responses stored in the local queue prior to memory responses stored in the buffered queue.

28. The memory system of claim 27 wherein the arbitration control logic develops the control signal to output memory responses stored in the buffered queue prior to memory responses stored in the local queue.

29. The memory system of claim 26 wherein the arbitration control logic develops the control signal to provide the memory responses from the local queue, buffered queue, and bypass path as a function of the age of a memory request associated with each memory response.

30. The memory system of claim 29 wherein the arbitration control logic assigns a time stamp to each memory request when the request is received by the hub, and wherein the age of each request corresponds to the assigned time stamp.

31. The memory system of claim 30 wherein the arbitration control logic develops the control signal to alternately output a predetermined number of memory responses stored in the buffered queue and a predetermined number of memory responses stored in the local queue.

32. The memory system of claim 31 wherein the same predetermined number of responses are alternately output from each queue if such predetermined number or a greater number of responses are stored in each queue.

33. The memory system of claim 24 wherein each of the local and downstream memory responses comprise data and a header identifying a memory request corresponding to the memory response.

34. A computer system, comprising:

a processor;

a system controller coupled to the processor, the system controller including a memory hub controller;

an input device coupled to the processor through the system controller;

an output device coupled to the processor through the system controller;

a storage device coupled to the processor through the system controller;

a plurality of memory modules, each memory module being coupled to adjacent memory modules through respective high-speed links, at least one of the memory modules being coupled to the memory hub controller through a respective high-speed link, and each memory module comprising:

a plurality of memory devices; and

a memory hub coupled to the memory devices and coupled to the corresponding high-speed links, the memory hub including,

a local queue adapted to receive local memory responses, and operable to store the local memory responses;

a bypass path adapted to receive downstream memory responses, and operable to pass the downstream memory responses;

a buffered queue coupled to the bypass path and operable to store downstream memory responses;

a multiplexer coupled to the local queue, the buffered queue and the bypass path, and operable to output responses from a selected one of the queues or the bypass path responsive to a control signal; and

arbitration control logic coupled to the multiplexer, the arbitration logic operable to develop the control signal to control the selection of responses output by the multiplexer.

35. The computer system of claim 34 wherein each of the high-speed links comprises an optical communications link.

36. The computer system of claim 34 wherein at least some of the memory devices comprise SDRAMs.

37. The computer system of claim 34 wherein the processor comprises a central processing unit (CPU).

38. The computer system of claim 34 wherein each of the local and downstream memory responses comprise data and a header identifying a memory request corresponding to the memory response.

39. A method of processing and forwarding memory responses in a memory system including a plurality of memory modules, each memory module including a memory hub coupled to memory devices, and the method comprising:

storing local memory responses from the memory devices;

storing downstream memory responses from downstream memory modules;

applying in each hub an independent arbitration algorithm to determine an order in which the stored memory responses are forwarded to an upstream memory module; and

forwarding the memory responses upstream according to the determined order.

40. The method of claim 39 wherein stored local memory responses are forwarded upstream prior to forwarding downstream memory responses upstream.

41. The method of claim 39 wherein stored downstream memory responses are forwarded upstream prior to forwarding local memory responses upstream.

42. The method of claim 39 wherein stored local and downstream memory responses are alternately forwarded upstream.

43. The method of claim 42 wherein a first number of local memory responses are forwarded upstream and a second number of downstream memory responses are forwarded upstream.

44. The method of claim 43 wherein the first and second numbers are equal.

45. The method of claim 39 wherein the stored local and downstream memory responses are forwarded upstream as a function of an age of a memory request associated with each memory response.

46. The method of claim 45 wherein a time stamp is assigned to each memory request when the request is received by each hub, and wherein the age of each memory request corresponds to the assigned time stamp.

47. The method of claim 39 wherein each of the local and downstream memory responses comprise data and a header identifying a memory request corresponding to the memory response.